

**IN THE UNITED STATES PATENT AND TRADEMARK OFFICE
(Attorney Docket No. 15550US02)**

<i>In the Reissue Application of:</i> John T. Bretscher	<i>Examiner:</i> Ajay M. Bhatia
<i>Application No.:</i> 10/821,833	<i>Group Art Unit:</i> 2145
<i>Filed:</i> April 9, 2004	<i>Confirmation No.:</i> 1417
<i>For:</i> COMPUTER SYSTEM ARCHITECTURE AND METHOD FOR MULTI-USER, REAL-TIME APPLICATIONS	<i>Express Mail Label No.</i>
<i>Original Patent:</i> 6,370,564	<i>Date of Express Mailing</i>

**DECLARATION OF PRIOR INVENTION IN THE UNITED STATES
TO OVERCOME CITED PATENT OR PUBLICATION (37 C.F.R. 1.131)**

Sir:

1. This declaration is to establish completion of invention in this application in the United States at a date prior to 21 March 1996, the effective date of U.S. Patent No. 5,828,843 (Grimm et al.) that was cited by the Examiner in the preceding Office Action.
2. The person making this declaration is the inventor named in the above-identified patent application.
3. I have an interest in a corporation which is in a corporate family that owns the above-referenced patent application as part of a portfolio, and I would receive some compensation attributable profit from the portfolio.
4. Attached, and incorporated by reference, is the Invention Data Disclosure Sheet ("IDDS") that I submitted to Ameritech that is dated August 20, 1995. Attention is drawn there to, including:
 - A. the drawings.
 - B. page 26: "No test of the invention has yet been made but Ameritech is building a fully functional prototype that may be operational by the end of August, 1995."
 - C. page 27: development began on June 27, 1995.
 - D. page 27: the problem was identified in January, 1995.
 - E. footnote 10 on page 25, the IDDS says that a single-board computer is one without its own hard disk, display, or keyboard
5. Conception of an embodiment was orally disclosed to another (Ali Shadman, another Ameritech employee) on June 9, 1995.

6. The first drawing was made on June 19, 1995, and the IDDS is the first written description.
7. The first disclosure to a person outside of Ameritech was made to Jim Hanson of Viacom New Media, on June 30, 1995, under an NDA.
8. Thereafter, work was carried out to complete prototypes, and the reduction to practice that is set out in the patent application at Col. 8, lines 31-67, occurred at a time prior to 21 March 1996, as follows:

A prototype of the computer architecture shown in FIG. 7 has been built, representing a preferred embodiment of the invention. The user stations 102, of course, will vary depending on the type of equipment chosen by the user. In the prototype built, one user station 102 was a personal computer based on the Intel Pentium processor, operating at 100 MHz with 32 MBytes of RAM and an 800 MByte hard disk. This user station 102 was also equipped with a quad-speed CD-ROM drive, a 32-bit sound card, and a telephone headset to free the user's hands while speaking into the telephone. A second user station 102 in the prototype was a personal computer based on the Intel 80486 processor, operating at 66 MHz with 4 MBytes of RAM, a 400 MByte hard disk, a dual-speed CD-ROM drive, a 16-bit sound card, and a telephone headset. Both user stations 102 were equipped with Multi-Tech model MT2834PCS multimode modems from Multi-Tech of Mounds View, Minn.

The prototype used a Hewlett-Packard model 735/125 workstation running the HP-UX operating system for the front-end server 112, and for the terminal server 110. The modem pool 108 employed Multi-Tech model MT2834PCS multimode modems, and was connected to the terminal server 110 using a serial interface and a RS-232C connection. In the prototype, the LAN 118 coupling the front-end server 112 and the plurality of dedicated processors 116 was an Ethernet LAN.

The voice bridge 120 in the prototype was a compute platform with an MVIP bus running the UNIX operating system. An 8-port analog voice board (AG-8 from Natural MicroSystems of Natick, Mass.) and an XDS/MVIP conference board from Amtelco of McFarland, Wis. provided the program-controlled teleconference functions. The interface between the modem pool 108 and the voice bridge 120 was a standard two-wire telephone connection. A LAN interface card facilitated the connection between the voice bridge 120 and the front-end server 112 via the LAN 118.

9. To the extent that the foregoing is not sufficient, work was also done in preparing and filing the original patent application, which led to two patents (U.S. Patent No. 6,370,564 and U.S. Patent No. 6,175,854) and the above-identified reissue patent application.

10. I worked on the patent application and cooperated with patent counsel at the law firm of Brinks, Hofer, Gilson, & Lione so that it could complete and file the original patent application. Consistent with that time line, I reviewed drafts of the patent application and provided comments to the patent counsel; I met with patent counsel to carefully go over a draft of the patent application; I also provided text additions to patent counsel; I approved the final version of the original patent application via communicating through Ameritech's in-house patent attorney.

11. I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

Respectfully submitted,

Date: October 8, 2008

John Bretscher
John Bretscher

Invention Data Disclosure Sheet

Privileged and Confidential
Attorney Work Product and
Attorney-Client Communication

Submitted by:
John Bretscher
2000 West Ameritech Center Drive
2c65a
Hoffman Estates, Illinois 60196
(708)248-5456

Invention Data Disclosure Sheet

Please type or print in ink. Attach additional sheets if necessary.

Please answer the following questions in as much detail as possible. Use drawings and equations where necessary.

1. What is the title or subject of the invention?

The invention is an architecture for a service that provides multiple-user real-time collaborative applications supporting voice and data. Each application instance runs on a compute server dedicated to that instance for the lifetime of the instance so that there is no contention and real-time response is achievable. The initial applications envisioned are multi-player, real-time response games.

2. Identify everyone who participated in the development and testing of this invention, and for each such individual describe the role that they played:

(a) Name: John Bretscher
Home Address: 515 Division Street
City/State: Elgin/Illinois
Citizen of: United States of America
Role Played: Inventor

3. What is the general product or process to which this invention relates? What is the specific product or process to which this invention relates?

This invention relates to multiple-user computer applications and, more specifically, to service centers providing computing resources for such applications where real-time response is important and where the users are geographically dispersed. The applications may be, but need not be, computer games.

4. Describe in detail (attach drawings if available) the current products or processes used by your company, and by others in the industry (i.e., how were things done before this invention).

The currently used methods are presented below in a logical order: each description is made in view of the ones previous and assumes that the previous ones have been read and understood. As the most pertinent application, I will in this section often refer to computer games for illustration. While using games as a mental reference point, remember that the invention is not restricted to game applications.

Monolithic Server: Local Access Method

Physical Description

One central computer plays host to numerous user stations, locally connected. Each user station may be a terminal or a computer emulating a terminal. It may even be a simple game input device (e.g., joystick) with no display at all.

Functional Description

All processing is performed on the central computer. The user stations may all display the same scene or each may display a scene tailored to the local user.

A popular application of this Method is when two users connect their game input devices into a shared personal computer (PC) or dedicated game device. The single display is then physically part of the server but is logically part of both user stations.

Advantages

Dedicated game devices are much cheaper than PCs so this is an easy way to start playing games. Because of this, there is a large base of these devices and of compatible games.

Since each user has a dedicated connection, access line latency is both low and fixed.


If the central computer is a dedicated game device, then there are no problems with competing applications slowing down play.

Since all processing is performed on one machine, there are no issues with event resolution ("refereeing").

If the users are in the same room, then they can freely talk with each other.

Heritage

Monolithic Server: Local Access Method



Modem-to-Modem Method

Physical Description

Two (and only two) user stations are connected by way of a modem link. Each user station has everything needed to support a single user in a stand-alone application, in addition to the ability to drive the modem. Usually the user station is a PC.

Functional Description

Each user station does all local processing and performs all display functions.

Usually one station (the "master") is chosen to resolve event contentions: to do this it receives inputs from both the local user and from the remote one, resolves contention between these inputs thus deciding the temporal progression of the application, and then communicates status of that progression. The master may tell each station only what it needs to know or may send the complete status to both stations and let each decide how much of that status is relevant to its local display.

There are also ways to share the event contention resolution function between the stations.

Advantages

Although PCs are expensive compared with dedicated game servers, their purchase is often justified by the fact that they have some use beyond games. Many homes already have a PC and a modem so there is little incremental cost involved beyond the application itself and the optional game input device. Also, having the enormous processing power of a PC dedicated to one user allows for graphics that are usually beyond the capability of a monolithic server that has to run everything.

The bandwidth and latency¹ of the dial-in line are fixed, unwavering, and the applications developers can often use this feature to effectively mask the fact that the bandwidth is fixed at a fairly low value (compared with local access) and the latency is fixed at a fairly high value (again compared with local access).

Each user controls his own station environment so there should never be contention from competing applications.


¹Bandwidth is a measure of the rate at which data can be sent over a communications link. It is usually measured in bits per second. Latency in a communications link measures the time from when a data bit enters the link at the data source to when it leaves the link at the data sink. It is often measured in milliseconds. Bandwidth and latency may vary completely independently of each other.

Two people from anywhere in the world can work together and, once the multimode modem² standards are complete, can talk with each other while they work.

²This is a new term that may or may not become a standard one. It refers to modems that can transmit voice at the same time and over the same link that they transmit data. If multimode modems are used in the Modem-to-Modem Method, then the application communicates over the data channel while at the same time, the voice channel lets the users talk to each other. They are also called Simultaneous Voice and Data (SVD) modems.

Heritage

Modem-to-Modem Method



Local Area Network Method

Physical Description

Several user stations (the number is limited by the specific application: usually it is four or eight) are connected on a Local Area Network (LAN). The user stations are similar in concept to those in the Modem-to-Modem Method with the obvious difference that LAN support replaces modem support.

Functional Description

The split between local processing and remote seen in the Modem-to-Modem Method also applies here. Usually, event resolution is done the same way with the master now resolving four to eight users instead of only two. Otherwise, this function can be handled by a separate server on the LAN.

Advantages

Most advantages of the Modem-to-Modem Method apply here also. The LAN's greatest additional advantage is that it allows more than two users in a single application.


As long as the LAN is not overloaded with traffic from other applications, it offers the high bandwidth and low latency that applications developers crave. Indeed, the LAN has so much bandwidth that it can be used for software distribution. Also, the LAN's broadcast mode can be used to reduce overall traffic requirements.

The processing burden placed upon a user station's CPU by the access line is reduced because LAN cards have their own communications processors.

If the users are physically proximate, then they can freely talk with one another.

Heritage

Local Area Network Method



Monolithic Server: Packet Network Method

Physical Description

This is the typical Internet game scenario. Several PC user stations (the number is limited by the application, as in the LAN Method) connect to a packet network. There are a few ways to do this but usually the connection is a modem link to an access server on the network. This access server may, or may not, be the same machine as the monolithic compute server. If it is the same, then this case degenerates to the Monolithic Server: Modem Access Method (which see).

For this section, I will assume that the access server and the compute server are distinct. They communicate through a shared packet network.

Functional Description

There are two common application schemes supported here. In the distributed scheme, local processing is done in the same way as in the LAN Method. The compute server is the master, doing all event contention resolution.

In the centralized scheme, all processing is done by the compute server and the user stations are mostly dumb input and output devices, just as in the Monolithic Server: Local Access Method.

Advantages

Some advantages of PC-based applications were described in the Modem-to-Modem Method. The number of people who subscribe to on-line packet network services is growing rapidly and broadening their usage habits to include multi-user applications is a natural move.


Two advantages accrue to the service provider: first, centralized applications need no software distribution channel. This makes offering new applications both easy and less of a financial risk. Second, connection to the shared packet network means that the compute server can be placed anywhere that is convenient for the service provider.

As an advantage over the LAN Method (where there is usually no central server), no user station is unfairly shackled with the burden of event contention resolution.

As an advantage over the Modem-to-Modem Method and the LAN Method, more than two people, placed anywhere the Internet reaches, can work together.

Heritage

Monolithic Server: Packet Network Method



Monolithic Server: Modem Access Method

Physical Description

As described above, this can be thought of as a degenerate case of the Monolithic Server: Packet Network Method. Everything there applies here with but few exceptions.

Functional Description

See Monolithic Server: Packet Network Method above.


Advantages

This adds the advantages of a dedicated access link, described in the Modem-to-Modem Method, to the advantages of the Monolithic Server: Packet Network Method. This is no small advantage because many applications, especially "high-twitch" games³, do not fare well under conditions of variable latency. This Method, quite possibly not implemented by anyone, is pretty close to what we want to do.

³These are games that demand very fast response from the players, so fast that the responses of a good player seem more reflexive than calculated. The games are often "shoot 'em ups" but need not be.

Heritage

Monolithic Server: Modem Access Method



5. What are the disadvantages of using each of these current products or processes?

While each Method introduces its own set of distinct disadvantages, most disadvantages are shared among several Methods. Rather than typing the same thing repeatedly, I assign letters to the disadvantages and refer to them that way. These reference letters are also used in the answers to questions 6, 7, 8, and 9.

Monolithic Server: Local Access Method

- [a] Users must be located close to each other.
- [b] Users may still be far enough apart that they cannot talk to each other.
- [c] The number of users is limited (often to two with dedicated game devices).
- [d] If the central server is not dedicated to one application instance, applications may compete for processing cycles and affect each other's performance.
- [e] One processor does everything so its performance sets limits on the kinds of applications that can be supported. In particular, support for graphics suffers.
- [f] Monolithic servers are not flexible in size or through time: they cannot adapt readily to accommodate changing demand and once purchased, quickly become obsolete.

Modem-to-Modem Method

- [c] (The applications are strictly limited to two users.)
- [g] PCs and communications equipment are expensive when compared with dedicated game devices.
- [h] Compared with the LAN Method, each user station must shoulder a greater processing burden in running its access line.
- [i] Compared with some other Methods, the bandwidth of the connection is low and the latency may be high.
- [j] If event contention resolution is not shared, the master must spend some of its processing power performing this function, possibly to the detriment of the local user.
- [k] If event contention resolution is shared, extra latency is usually added.

Local Area Network Method

[a], [b], [g]

- [l] Few people have access to a LAN, and then often only at work. In any case, it is difficult to gather together enough people to make collaborative applications worthwhile (except for applications that support people at work doing work, of course).
- [m] Latency and bandwidth are both heavily dependent upon the characteristics of traffic on the LAN, which may be subject to enormous variation.
- [n] The event contention resolution master must deal with more users than in the Modem-to-Modem Method, possibly to the detriment of the local user.

Monolithic Server: Packet Network Method

[d], [f], [g], [h]

- [o] Latency and bandwidth are both heavily dependent upon the characteristics of traffic on the packet network, which may be subject to enormous variation. This reason parallels reason [m].
- [p] For centralized applications, one processor does everything so its performance sets limits on the kinds of applications that can be supported. This is similar to reason [e].
- [q] Depending upon the bandwidth of the access link, there may need to be an outside distribution channel for distributed applications.
- [r] In almost all cases, the access link runs over the user's telephone line. Given that, there seems to be no good way to add a voice capability.

Monolithic Server: Modem Access Method

[d], [f], [g], [h], [i], [p], [q]

- [s] Compared with the Monolithic Server: Packet Network Method, the economics of the service are much more dependent upon the relationship between the server sites and the incoming user traffic. It is more difficult to economically distribute the service over a wide area.

6. Are there any ways to minimize the adverse effects of these disadvantages?

In the answer to question 4, I presented five different Methods, each of which has certain advantages and certain disadvantages. The obvious way to minimize the adverse effects of almost any of these disadvantages is to use a different Method. In this section I will just assume that this is understood and only list ways that apply within a given Method. The bracketed letters refer to the disadvantages described in the answer to question 5.

- [a] Depending upon the specific access technology, there may be ways to extend the distance. For example, LAN bridges can extend the presence of a LAN almost indefinitely.
- [b] Often the telephone network can be used. For the LAN Method, also see the comments on [m] and [o] below.
- [c], [e] For a monolithic server, the number of users may be limited by the server's processing capability or by the number of ports it can support. More powerful servers could alleviate this problem for a while. The limit of two users in the Modem-to-Modem Method cannot be addressed within the Method.
- [d] A server can be dedicated to just one application instance.
- [f] Multiple-processor server machines may address this. (Our solution is similar to this one.)
- [g] Lease options for the equipment could be provided.
- [h] A more powerful user station processor or a greater serial port data rate can make this disadvantage much less apparent.
- [i] The highest speed modems still cannot compete with some other Methods. Modems may be specially built that minimize latency at the expense of bandwidth. This problem can be minimized by applications developers who design with the limits in mind.
- [j], [k] Each application must strike a balance between processor response and latency requirements.
- [l] Remote LAN bridges may bring more people together. This is similar to the comment in [a], above.

- [m], [o] Isochronous⁴ data channel standards are coming for many LANs and packet networks. These channels can be used for data or for voice.
- [n] The event contention resolution master may be dedicated to that task and have no local user.
- [p] Applications can be built to the distributed scheme.
- [q] A faster or a cheaper access link can be used.
- [r] Users may have a second telephone line for voice. Or they may wait for the packet isochronous standards, as in [m] and [o] above.
- [s] Special trunking arrangements can be set up that decrease the transport costs across the telephone network. Also, marketing studies can indicate where best to place service centers and how big such centers should be.

⁴An isochronous channel is one whose latency never varies. People would have a very hard time understanding standard telephone voice calls if those connections were not isochronous.

7. Are there other disadvantages, such as prohibitive cost or labor intensity, in attempting to minimize these disadvantages?
- [a], [l] This may be quite costly on a per user basis. Security risks are posed that may be unacceptable if there are corporate resources involved. Administrative overhead must be borne by someone.
 - [b], [r] Tying up a telephone line for the duration of the application can be both expensive and hard to justify. A second telephone line is an ongoing expense. Bridging voices (for more than two users) is an added expense and complication. These costs apply to each user individually, rather than once for the service provider. Also see [m] and [o] below.
 - [c], [e], [f] No matter how powerful the server is initially, the demands of applications and users will keep growing until the server can no longer keep up, then it becomes obsolete. Huge servers are both very expensive and not very cost effective. They are not standardized so the service provider will be tied into one supplier and one architecture. Significant administrative overhead must be borne by the service provider.
 - [d], [n] A dedicated server is an additional cost and one may not be available. Since it has no local user, someone must step up to the task of administering it.
 - [g] Leasing is expensive for the user and that cost applies to each user individually. Finding someone willing to support a lease on equipment with a short life span will be difficult.
 - [h] This is just money but the cost applies to each user individually. Also, there is already a large installed base of this equipment and users may be reluctant to change.
 - [i] Special modems are an administrative nightmare for everyone involved. The demand for the service may not initially be great enough to increase demand for the modems and thus drive down their price. These special modems may not be very good for other applications (unless they use a Digital Signal Processor⁵ with download capability). There is a large installed base of modems. The cost applies to each user individually.

⁵Some modem manufacturers are putting Digital Signal Processors (DSPs) in their modems to improve performance under varying line conditions. These DSPs are often remotely programmable which allows the manufacturer to upgrade the modem without a hardware swap out. Could not the manufacturers use this capability to alter the fundamental processing of the modem for special applications (such as ours), always resetting the modem to its default state later?

- [j], [k] The compromise may eliminate the worst aspects of the disadvantages but it leaves both processor response and latency at suboptimal levels.
- [m], [o] The isochronous standards are not yet complete. There is a large installed base that does not meet the new standards. Every user must change over to the new standards.
- [p] Setting up a software distribution channel is costly. Version control can be a problem for the service provider, the applications developers, and the users.
- [q] This is expensive and the cost applies to each user individually. Faster links are not available in many places. The service provider's administration budget must increase.
- [s] This is difficult to do correctly and is an ongoing cost for the service provider.

8. Does your invention attempt to overcome each of the disadvantages identified in Question 5? Does it overcome those disadvantages?

Since the "how" and the "why" are relegated to question 9, I here give just the "whether".

The invention, properly understood, is but one part of a complete architecture that supports the service. It is, however, a part which enables all other parts to function optimally. In the answer to the current question, I refer to the capabilities of the service architecture as a whole. In the answer to question 9, I am much more explicit and specific.

	Attempt?	Succeed?
[a]	Yes	Yes
[b]	Yes	Yes
[c]	Yes	Yes
[d]	Yes	Yes
[e]	Yes	Yes
[f]	Yes	Yes
[g]	Yes	Partially
[h]	No	No
[i]	Yes	Partially
[j]	Yes	Yes
[k]	(This does not apply.)	
[l]	Yes	Yes
[m]	Yes	Yes
[n]	Yes	Yes
[o]	Yes	Yes
[p]	Yes	Yes
[q]	Yes	Partially
[r]	Yes	Yes
[s]	Yes	Partially

Are there any other purposes that your invention seeks to achieve?

Generally, the invention seeks to bring together the advantages of each of the five Methods described in the answer to question 4 while partitioning off any remaining disadvantages so that they can be dealt with most effectively. While there are many other purposes of the service as a whole, I can not think of any others for the invention itself, narrowly understood.

9. Describe in detail how your invention overcomes those disadvantages (attach additional sheets, drawings, sketches, etc., if needed). Be sure to explain why your proposed solution is better than other proposed solutions.

The distinction raised in the answer to question 8 between the invention, narrowly understood, and the service architecture as a whole is important to maintain in the answer to this question. I preface each answer with an "I" if the invention provides this advantage and an "A" if the service architecture as a whole is responsible. As noted above, the invention enables the service architecture but is not responsible for all of its many advantages.

- [a] A: Users may be located anywhere there are telephones. (Actually, satellite links would probably provide an unacceptable level of service).
- [b], [r] A: Users can talk among themselves over the voice channels provided by the multimode modems. Thus there is no need for a second telephone line.
- [c] A & I: The number of simultaneous users of an application is limited only by the design of that application. Since a separate compute server is dedicated to each application instance, this limit may be very high.
- [d], [n] I: A compute server is dedicated to each application instance. It has no local user to support and is also pretty inexpensive. This dedication gives the applications designers enormous freedom and power.
- [e], [p] A & I: Since graphics processing (and the like) can be performed locally at each user station, the compute server is freed up to do only what it alone can do. As in [d] and [n] above, the applications designers benefit from this.
- [f] I: There are no monolithic servers. The pool of compute servers can easily change to meet changing demand or obsolescence. The compute servers are also standard items, giving the service supplier some vendor freedom.
- [g] A: The service provider could provide modem leases. Users could use regular modems if they are willing to give up the voice capability.
- [h] This disadvantage is not addressed.
- [i] A: The bandwidth and latency are fixed and applications designers can optimize to these conditions. The architecture helps to minimize the amount of data that must be sent so that the limitations are not as important.

- [j], [k] I: Event contention resolution is not shared: it is performed by the compute server. Since the compute server does not need to support a local user, there is no need to compromise between processor response and latency.
- [l] A: Since users can dial into the service, there are no technical problems with getting enough of them.
- [m], [o] A: Latency and bandwidth are fixed on the dial-in lines and isochronous channels will be implemented within the server complex. Only the service provider needs to adapt to these forthcoming standards, not the individual users.
- [q] A: Although the disadvantages of an outside distribution channel are not directly addressed, at least the server complex could check versions and alert users as needed.
- [s] A & I: Ameritech is very good at this kind of thing. By design, the invention is quite transportable if market needs ever change.

Here is a simple summary listing some of the major advantages of this service architecture, which incorporates the invention, over the five methods described in the answer to question 4.

This service architecture is better than the Monolithic Server: Local Access Method because it allows users to be geographically dispersed while still in voice contact, it allows for more users in an application, and it leverages the power of individual user stations.

This service architecture is better than the Modem-to-Modem Method because it allows for more than two users in an application.

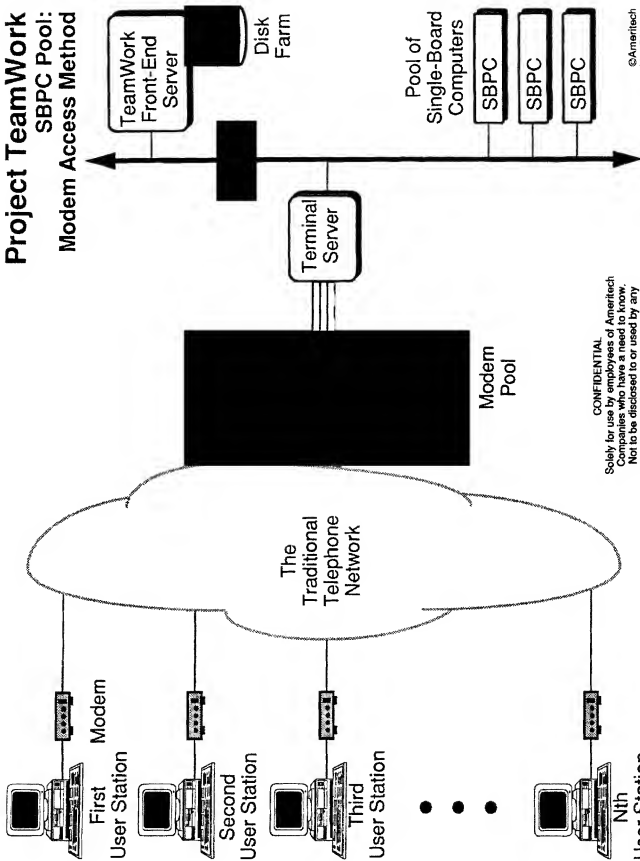
This service architecture is better than the Local Area Network Method because it allows the users to be geographically dispersed while still in voice contact.

This service architecture is better than the Monolithic Server: Packet Network Method because its isochronous channels allow for both voice and real-time response and because of its pool of compute servers (and all that in it inheres).

This service architecture is better than the Monolithic Server: Modem Access Method because of its pool of compute servers.

Project TeamWork

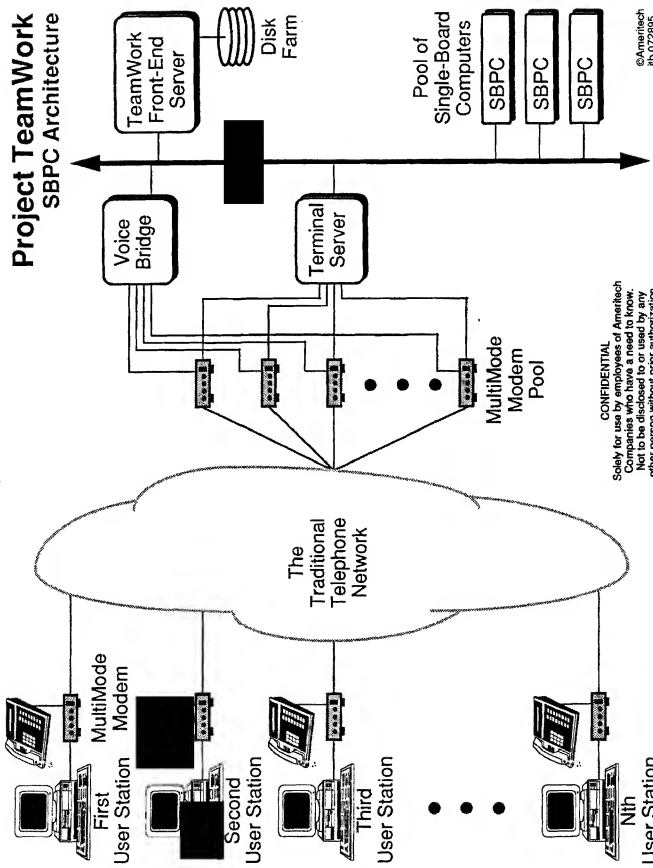
SBPC Pool: Modem Access Method



CONFIDENTIAL
Solely for use by employees of Ameritech
Companies who have a need to know.
Not to be disclosed to or used by any
other person without prior authorization.

©Ameritech
jfb 072895

Project TeamWork SBPC Architecture



CONFIDENTIAL
Solely for use by employees of Ameritech
Companies who have a need to know.
Not to be disclosed to or used by any
other person without prior authorization.

10. What is the best way to practice the invention? What are the particular sizes, shapes or dimensions of a new product? What are the times, temperatures and pressures that are known to give the best results if this is a new chemical process? If successful tests of the invention have been made, identify and describe them.

One real advantage of this invention is that it does not force a choice of a particular method of practice. While we do not yet know the best way to practice the invention, I will share thoughts on some of the good ways.

The method used to access the server complex, while not part of the invention itself, is vital to successful functioning. That method should provide an isochronous, full-duplex⁶ data channel between each user and the compute server in order to take full advantage of the invention's ability to support real-time response applications. The traditional telephone network provides such channels. So does ISDN⁷. Packet access technologies may provide them in the future.

Similarly, the data channels must be isochronous within the server complex itself and some LAN technologies do not provide for that. We are investigating Token Ring, FDDI⁸, Ethernet 100BaseT, switched Ethernet, and ATM⁹ LANs for deployment.

Initially we are looking only at PCs for the servers and for the user stations but that is just to make the project start-up easier: this is not a philosophical choice nor a long term limitation.

Ideally, each Single Board Personal Computer (SBPC)¹⁰ runs only one application instance at a time. This will make the application developer's job much easier.

The invention does not depend upon having a separate voice channel to each user but neither does it prevent it. Our opinion is that voice makes the service much more compelling. We will initially offer both voice and non-voice levels of service.

⁶A full-duplex channel between A and B can be thought of as having two subchannels, one in the "forward" direction from A to B and one in the "reverse" direction, from B to A. A and B can each always transmit, regardless of what the other one is doing. The two do not effect each other at all.

⁷The Integrated Services Digital Network is an access technology that usually gives an end user two independent digital channels running at 64 Kbps each, suitable for voice or data, and one digital channel running at 16 Kbps, suitable for data and signaling.

⁸The Fiber Distributed Data Interface is a standard LAN technology. Running at 100 Mbps, it is in essence a fast Token Ring.

⁹Asynchronous Transfer Mode is a slimmed down network protocol that happens to be very useful for implementing high-speed LANs. It does not specify an access rate but speeds generally start at 56 Mbps and go on up into the Gbps range.

¹⁰A Single Board Computer may consist of a processor, Random Access Memory, Read-Only Memory, a LAN interface, and a power supply. By leaving out the hard disk, the display, and the keyboard, the Single Board becomes quite inexpensive and is well suited to the narrow task of processing data (which may mean running a game stored somewhere else).

That said, the best way to handle voice may be to move the voice bridge function onto each SBPC. We cannot do that initially but in the future it will give the application developers full control over voice as a sound effect.

No test of the invention has yet been made but Ameritech is building a fully functional prototype that may be operational by the end of August, 1995. The prototype will be used for, among other things, testing the concepts of the invention. Please see the attached documents "TeamWork Issues and Concepts: Prototypes" and "SBPC Issues and Concepts: Prototype".

11. Are there other ways to practice the invention? Usually, it is possible to vary the operation, structure, composition, or other aspects of the invention without losing its advantages. Be sure to identify those aspects that can be varied without losing the advantages of the invention.

The invention does not stipulate an access method. Although the diagrams show users connecting via the switched telephone network, the connections could also be ISDN, packet network, Hybrid Fiber Coax, LAN, or whatever.

The invention is also independent of the method of connecting the various pieces of the server complex. A LAN is one obvious way to connect but there are alternatives.

The bank of compute servers need not be homogenous: any possible machine could fit in. Also the compute servers could have local disks or any other peripherals desired. The point of the invention is that no compute server needs to divide its time between multiple real-time applications.

Although the invention is focused on multiple-user, real-time applications, it certainly can and will support numerous applications that do not fit that definition.

Keeping in mind the two immediately preceding paragraphs, each SBPC may run more than one application instance at one time. I do not believe that any reduced cost would ever offset the added complication and the application developer's effort, unless the applications do not require real-time response.

This invention, strictly understood, has nothing directly to do with the user voice channels. The invention puts up no barriers to their integration, but they are not strictly required. As such, the voice bridge function can be put anywhere.

12. Please provide dates and supporting documents for each of the items listed below:

In this section, I refer to the invention as strictly defined.

Conception of the Invention

When was the problem identified?	January, 1995
First oral disclosure:	Date: June 9, 1995
	To whom: Ali Shadman
First drawing:	Date: June 19, 1995
	Drawing number: <none>
First written description:	Date: This is it.

Development of the Invention

Date work on development began:	June 27, 1995
Date development work completed:	No time soon.
Date of first prototype or model:	Planned for end of August, 1995
Date of first successful test:	Planned for end of August, 1995

First Disclosure Outside the Company

Has the discovery been disclosed to anyone outside the company or published in any manner? If so, by whom, when and where?

Yes, the invention was presented by John Bretscher (and assorted other Ameritech folk) to Jim Hanson of Viacom New Media during a meeting on June 30, 1995 in the Ameritech Center. Ameritech has a non-disclosure filed with Viacom. There has been no publication.

First Commercial Use or Sale

Has the discovery been shown, given, advertised, offered or sold to anyone outside the company? If so, by whom, when, and to whom?

No (other than shown to Jim Hanson as described above).

13. Identify the most closely related prior publications, prior patents and prior products or uses.

Diane Peterson of Ameritech sent you a stack of related prior publications. Possible prior products include DWANGO, Catapult, and ICTV.

14. Please sign below:

The foregoing invention disclosure consisting of 28 pages was read and understood by me on the date opposite my name.

Contributor:	<u>John Bretscher</u>	Date:	<u>August 25, 1995</u>
Reviewers:	(1) <u>Bonglishing</u>	Date:	<u>8/25/95</u>
	(2) <u>Dale H</u>	Date:	<u>8/25/95</u>
Witnesses:	(1) <u>Jim Buckley</u>	Date:	<u>8/29/95</u>
	(2) <u>[Signature]</u>	Date:	<u>8/30/95</u>

Prototypes address several issues: they may test how existing components work together in new ways; they may test new hardware and software in an environment close to the deployment target; they may be used to measure customer perceptions of individual aspects of the platform or of the service as a whole; they can point to problems that will only completely emerge during full-scale deployment; and they can be used to convince management of the desirability of proceeding with the project. It may be impossible to design a single prototype that addresses all of these issues, but it is important to TeamWork that each issue be addressed by one or by a combination of prototypes.

TeamWork has already displayed one concept prototype and is involved in building a prototype that is fully functional and that points the way to a service deployment architecture. I will describe these prototypes (and one other, not currently planned) by their relation to the issues mentioned above.

MultiMode Modem Concept Prototype

There are two orthogonal theories driving the TeamWork project. First, people that interact through computer applications may also like to talk to each other. Second, some real-time response applications may benefit from the ability to support more than two remote users.

To test the voice theory, we turned to the PC game market. There are several games currently available that allow two people, at a physical remove, to play together after connecting their PCs via modems. Without investing too much money or time, we wanted to see if adding a voice channel to this environment would significantly enhance the experience for the players.

We put together a simple prototype using widely available components and commercially available games. The only new part was the modems: we used multimode modems that allow, simultaneously, a voice and a data connection between the players. The games used the data connection just as they would on normal modems while the simultaneous voice channel allowed the players to talk to each other.

Several Ameritech people tried out numerous two-player games and we were able to conclude that the voice channel truly vitalizes the experience and changes the perceived context from playing with a computer to playing with another person, mediated by computers. This social perception shift may be the most important aspect of the entire gaming world to Ameritech.

Multi-User LAN Prototype

Since the concept prototype uses a modem connection, it is limited to only two users. To test the second theory, we need to extend our work to allow for multiple¹ users.

To do that using existing games, we could get a multi-player LAN game, set up a dedicated LAN, attach user stations in separate rooms, add a telephone to each room, and then connect the phones with a simple voice bridge. We would then have a good platform, set up for as many players as the game allows (usually four to eight) with full control of the voice connections. This could be used for demonstrations to internal management and for user response studies.

The major shortcoming of this prototype would be that since the connection medium is a LAN instead of modems, the response time may be quite a bit better for this prototype than for the actual TeamWork service. This presents the danger that user testing may provide overly optimistic results. Here are two suggestions about how to slow down the prototype:

- (1) Simulate a slower connection by choosing a slow PC as the event contention resolution master.
- (2) Bridge the LAN using ISDN. This adds extra complexity to the prototype but would slow it down.

Even without either of these methods, I think the advantages of low cost, easy transport, and quick implementation with no development make this attractive for a traveling prototype.

There are no plans at present to pursue this option.

SBPC Architectural Prototype

The service plan for project TeamWork calls for connecting multiple users together who, initially at least, will use modems as their access method. To allow for that, a special service platform must be built. The SBPC architecture is one proposal for how that platform should be built.

¹In this document, "multiple" or "multi-" user or player always refers to more than two. This is an important distinction since, as shown by the concept prototype, games limited to two players are already possible and do not need the TeamWork service platform.

August 20, 1995

CONFIDENTIAL

Solely for use by employees of Ameritech Companies who have a need to know. Not to be disclosed to or used by any other person without prior authorization.

We are currently putting together a full-blown SBPC architectural prototype that mimics in all important aspects² a deployable service platform. Up to four people will be able to dial into the prototype, socialize with the other users, and then actually run an application, all the while making use of the simultaneous voice and data connections. For technical details, please see "SBPC Issues and Concepts: Prototype".

Of course, this wonderful prototype remains less than compelling until there are some applications running on it. Ameritech does not envision itself writing those applications but Viacom New Media has volunteered to modify one of their games specifically for this prototype. Their offer is especially attractive to both parties because the prototype is so closely aligned with the service deployment platform. We will be working closely with them as the game and the prototype are debugged together.

Once everything is working, the prototype should be able to perform all of the functions listed at the start of this document. It will be ideal for testing customer reactions because it is not a mock-up of a proposed service, it really is the service, though on a smaller scale. From a logistics point of view, focus groups will be easy to deal with because the server complex can stay in an Ameritech lab while the user stations go out to meet the populace: it really is a dial-in service.

Though quite functional, this prototype can not address every conceivable issue. For example, service deployment issues such as operations systems and region-wide scaling are not addressed.

This prototype should be operational by the end of August, 1995.

²It is to be hoped.
August 20, 1995

The architectural prototype is meant to verify the feasibility of the SBPC concept for project TeamWork¹. As such, it is as close to an actual service deployment architecture as is possible. Here I describe what the prototype physically entails and specifically emphasize the choices made with an eye to service deployment. Since Viacom will be writing the first applications, they gave us valuable input into what kind of pieces would make their testing most relevant.

User Station

This prototype does not intend to certify applications over the entire range of possible user station configurations or even to set the boundaries of the desirable. But it is important, as a benchmark at least, to verify that disparate machines can work together in a reasonable way. The user station configurations below cover most of what we actually expect to see when the service is deployed. The prototype is flexible enough to add just about anything else if it ever becomes useful so to do.

Three Pentium PCs: probably 100 MHz or greater, PCI bus, 16 to 32 MB RAM, at least 800 MB hard disk, quad speed CD-ROM, Sound Blaster Pro or AWE 32 sound card, Plantronics SP-05 telephone headset, and some telephone base set, maybe with a speaker phone.

One 486 DX2 66 PC: VLB, VESA, four to eight MB RAM, at least 400 MB hard disk, double speed CD-ROM, Sound Blaster Pro sound card, Plantronics SP-05 telephone headset, and some telephone base set.

I do not yet know what kind of game input devices will be useful for testing. The applications we get will help determine our needs here.

MultiMode Modem

We can use this prototype to compare modems and, once the standards are done, test interoperability between modems. Initially we will use Multi-Tech MT2834PCS simultaneous voice and data modems.

Modem Pool

This will just be four modems identical in specification to the user station multimode modems.

¹For context, please also see "TeamWork Issues and Concepts: Prototypes".
August 20, 1995

Terminal Server

For the prototype, this will be the same Hewlett Packard workstation used for the front-end server (see below). The interface between the modems and the terminal server will be RS-232C. In this area, the prototype may depart significantly from the deployment architecture.

LAN

Ethernet is cheap and ubiquitous plus it has more than enough bandwidth:

Each user station link runs no faster than 28.8 Kbps minus at least 12 Kbps for voice leaving 16.8 Kbps for data. The connections are full duplex so we multiply that by two to get 33.6 Kbps per user station. There are four stations making the total station traffic 134.4 Kbps to which must be added overhead for Ethernet protocol headers and whatever other messaging is going on during game play. These later two figures are not known (indeed they will vary with the application and with the functionality of the front-end software) but since station traffic only accounts for 1.3% of the total bandwidth on an Ethernet, I do not see any problems for the prototype.

In this area, we are relying upon the small scale of the prototype to avoid doing the traffic engineering of the server complex LAN. Ethernet does not currently support the isochronous circuits that we will eventually need. This is going to be an important work area during the scale-up to deployment.

Voice Bridge

The voice bridge is a Unix box from Natural Micro Systems with an MVIP bus, 16 MB RAM, some hard disk, an AG8 board with 8 voice ports, and soon an Ethernet interface card. The interface between the modems and the voice bridge will be standard Tip and Ring.

Front-End Server

This is a Hewlett Packard Workstation, model 735/125MHzSPU with 144 MB RAM, and three hard disks: 3.6 GB, 1.6GB, 1.6 GB. It runs HP-UX. We are using it because we have it: we have made no study of front-ends.

SBPC Pool

Initially this will consist of just one Pentium PC (not a single board), configured in the same way as a user station. Later, as the functionality of our LAN operations code develops, we can remove functionality from this PC until we get it down to the level of a single board.

Non-Application Software

This is all the custom-developed software that supports and surrounds the running of the target applications. It falls readily into two heaps: the User Interface code and the LAN Operations code.

User Interface

We already recognize that the distinction between target applications and the user interface software may be meaningless to the actual users: the user interface, with its voice chat capability, may eventually become a more compelling experience than any application that runs on an SBPC.

But that's eventually. In the prototype, the user interface code has a somewhat more manageable set of goals: it gives the user a simple way to dial into the server complex; it runs some kind of chat room for users not yet involved in an application; it lets the users form ad hoc groups for the applications; it leads the user to the point where the application itself takes over; and it ushers the user out of the application or out of the service at the user's discretion.

LAN Operations

This is the real low-level stuff, the silicon plumbing, that ties once separate machines together into a working server complex and lets others concentrate on applications. Users should never be aware of any of this: even error conditions generated down here should only be presented to the user by means of the user interface code.

Briefly, the major functions that this code performs are: on SBPC reset, it waits for an application to be downloaded and started; it allows an application to be downloaded from the front-end server's disk farm to the SBPC's RAM; on a signal from the front-end, it starts the application with the users specified; it communicates user drops and maybe adds to the application; and it sends SBPC status to the front-end.

For much more detail about how this is all going to work, please see "SBPC Issues and Concepts: Message Flows".

Application Software

By application software, I mean the target applications, which is to say, games, at least in the short run. For internal testing of parts of the platform, we can use an existing two-player modem game running between one user station and the "SBPC". We may try to test with an existing, multiple-player LAN game but that may not work at all well.

Full service platform testing only begins when we have an application specifically written for this prototype. Viacom has volunteered to modify one of their games for this purpose. We should have it by the end of August.